

Privacy-Preserving Nonlinear Observer Design Using Contraction Analysis

Jerome Le Ny

Abstract—Real-time signal processing applications are increasingly focused on analyzing privacy-sensitive data obtained from individuals, and this data might need to be processed through model-based estimators to produce accurate statistics. Moreover, the models used in population dynamics studies, e.g., in epidemiology or sociology, are often necessarily nonlinear. This paper presents a design approach for nonlinear privacy-preserving model-based observers, relying on contraction analysis to give differential privacy guarantees to the individuals providing the input data. The approach is illustrated in two applications: estimation of edge formation probabilities in a dynamic social network, and syndromic surveillance relying on an epidemiological model.

I. INTRODUCTION

The development of many recent technological systems, such as location-based services, the “Internet of Things”, or electronic biosurveillance systems, relies on the analysis of personal data originating from generally privacy-sensitive participants. In many cases, the system is only interested in producing aggregate statistics from these individual data streams, e.g., a dynamic map showing road traffic conditions or an estimate of power consumption in a neighborhood, but even though aggregation helps, significant privacy breaches cannot be ruled out a priori [1]–[3]. This is mainly due to the possibility of correlating the system’s output with other publicly available data. The integration of privacy-preserving mechanisms with formal guarantees into such systems would help alleviate some of the justified concerns of the participants and encourage wider adoption.

While various information theoretic definitions can be given to the concept of privacy and are potentially applicable to the processing of data streams in real-time [4], we focus on the notion of differential privacy, which originates from the database and cryptography literature [5]. A differentially private mechanism publishes information about a dataset in a way that is not too sensitive to a single individual’s data. As a result, it becomes difficult to make inferences about that individual from the published output. Previous work on the design of linear filters with differential privacy guarantees includes [6]–[10]. The problem studied in this paper is that of designing privacy-preserving nonlinear model-based estimators, which to the best of our knowledge has not been studied in a general setting before.

A convenient way of achieving differential privacy for an estimator is to bound its so-called *sensitivity* [5], a form of

incremental system gain between the private input signal and the published output [9]. Various tools can be used for this purpose, and here we rely on contraction analysis, see, e.g., [11]–[14] and the references therein.

The rest of the paper is divided as follows. Section II presents the problem statement formally, provides a brief introduction to the notion of differential privacy, and describes privacy-preserving mechanisms with input and output perturbation. In Section III we develop a type of “vanishing-input vanishing-output” property of contracting systems similar to the one presented in [12] but stated here for discrete-time systems. This result is then applied in Section IV to the design of differentially private observers with output perturbation. The methodology is illustrated via two examples. In Section V, we consider the problem of estimating link formation probabilities in a dynamic social network, with a nonlinear measurement model. In Section VI, we consider a nonlinear epidemiological model and design a differentially private estimator of the proportion of susceptible and infectious people in a population, assuming a syndromic data source.

Notation: In this paper, $\mathbb{N} := \{0, 1, \dots\}$ denotes the set of non-negative integers. For $T : X \rightarrow Y$ a linear map between finite dimensional vector spaces X and Y equipped with the norms $\|\cdot\|_X$ and $\|\cdot\|_Y$ respectively, we denote by $\|T\|_{X,Y}$ its induced norm. If $X = Y$ and both spaces are equipped with the same norm $\|\cdot\|_X$, we simply write $\|\cdot\|_X$.

II. PROBLEM STATEMENT

A. Observer Design

Suppose that we can measure a discrete-time signal $\{y_k\}_{k \geq 0}$ for which we have a state-space model of the form

$$x_{k+1} = f_k(x_k) + w_k \quad (1)$$

$$y_k = g_k(x_k) + v_k, \quad (2)$$

where w_k, v_k are noise signals capturing the uncertainty in the model, $x_k \in X := \mathbb{R}^n$ for some n , and $y_k \in Y := \mathbb{R}^m$ for some m . The goal is to reconstruct from y_k an estimate of the state x_k that we denote z_k , i.e., we want to build a state observer, which we assume in this paper to be of the simple Luenberger-type form

$$z_{k+1} = f_k(z_k) + L_k(y_k - g_k(z_k)), \quad (3)$$

where L_k is a sequence of gain matrices to determine.

In the applications discussed later in the paper, the signal y_k is collected from privacy-sensitive individuals, hence needs to be protected. On the other hand, the model (1), (2), i.e., the functions f_k, g_k , is assumed to be publicly

This work was supported by NSERC under Grant RGPIN-435905-13. The author is with the department of Electrical Engineering, Polytechnique Montreal, and with GERAD, Montreal, QC H3T-1J4, Canada jerome.le-ny@polymtl.ca

available. The data aggregator wishes to release the signal z_k produced by (3) publicly as well. However, since z_k depends on the sensitive signal y_k , we will only allow the release of an approximate version of z_k carrying certain privacy guarantees detailed formally in the next subsection. We will later see that the gain matrices need to be carefully chosen to balance accuracy or speed of the observer on the one hand and the level of privacy offered on the other hand.

Remark 1: Note that we do not provide here nor use in our designs any model of the noise signals w_k and v_k , which are simply used as a device to explain the discrepancy between any measured signal y_k and the signal predicted by a deterministic model.

B. Differential Privacy

A differentially private version of the observer (3) should produce an output that is not too sensitive to certain variations associated to an individual's data in the input signal y_k . The formal definition of differential privacy is given in Definition 1 below. An individual's signal could correspond to a specific component of y_k , or y_k could already represent an signal aggregated from many individuals [9]. We specify first the type of variations in y_k that we want to make hard to detect by defining a symmetric binary relation, denoted Adj , on the space of datasets \mathcal{D} of interest, here the space of signals y . We consider here the following adjacency relation

$$\text{Adj}(y, \tilde{y}) \text{ iff} \quad (4)$$

$$\exists k_0 \geq 0 \text{ s.t. } \begin{cases} y_k = \tilde{y}_k, & k < k_0 \\ |y_k - \tilde{y}_k|_{\mathcal{Y}} \leq K\alpha^{k-k_0}, & k \geq k_0, \end{cases}$$

where $|\cdot|_{\mathcal{Y}}$ is a specified norm on \mathcal{Y} , and $K > 0$, $0 \leq \alpha < 1$ are given constants. In other words, we aim at providing differential privacy guarantees for transient deviations starting at any time k_0 that subsequently decrease geometrically. Note that in [6], [7] the authors consider for the design of a differentially private counter an adjacency condition where the (scalar) input signals can vary by at most one and at a single time period. In comparison, our adjacency condition (4) greatly enlarges the set of signal deviations associated to an individual for which we aim to provide guarantees.

Differentially private mechanisms necessarily randomize their outputs, so that they satisfy the following property.

Definition 1: Let \mathcal{D} be a space equipped with a symmetric binary relation denoted Adj , and let $(\mathcal{R}, \mathcal{M})$ be a measurable space. Let $\epsilon, \delta \geq 0$. A mechanism $M : \mathcal{D} \times \Omega \rightarrow \mathcal{R}$ is (ϵ, δ) -differentially private for Adj if for all $d, d' \in \mathcal{D}$ such that $\text{Adj}(d, d')$, we have

$$\mathbb{P}(M(d) \in S) \leq e^\epsilon \mathbb{P}(M(d') \in S) + \delta, \quad \forall S \in \mathcal{M}. \quad (5)$$

If $\delta = 0$, the mechanism is said to be ϵ -differentially private. This definition quantifies the allowed deviation for the output distribution of a differentially private mechanism, when the variations at the input satisfy the adjacency relation. Smaller values of ϵ and δ correspond to stronger privacy guarantees. In this paper, the space \mathcal{D} was defined as the space of input signals y , the adjacency relation considered is (4), and the

output space \mathcal{R} is the space of output signals z for the observer. We then wish to publish an accurate estimate of the state x while satisfying the property of Definition 1 for specified values of ϵ and δ .

C. Sensitivity and Basic Mechanisms

Enforcing differential privacy can be done by randomly perturbing the published output of a system, at the price of reducing its utility or quality. Hence, we are interested in evaluating as precisely as possible the amount of noise necessary to make a mechanism differentially private. For this purpose, the following quantity plays an important role.

Definition 2: Let p be a positive integer. The ℓ_p -sensitivity of a system G with m inputs and n outputs with respect to the adjacency relation Adj is defined by

$$\Delta_p G = \sup_{\text{Adj}(u, u')} \|Gu - Gu'\|_p$$

where by definition $\|v\|_p = (\sum_{k=0}^{\infty} \sum_{i=1}^n |v_{k,i}|^p)^{1/p}$ for $v = \{v_k\}_{k \geq 0}$ a vector-valued signal, where $v_k \in \mathbb{R}^n$ has components $\{v_{k,i}\}_{i=1}^n$.

In practice we will be interested in the sensitivity of a system for the cases $p = 1$ and $p = 2$. The basic mechanisms of Theorem 1 below (see [9] for proofs and references), can be used to produce differentially private signals. First, we need the following definitions. A zero-mean Laplace random variable with parameter b has the pdf $\exp(-|x|/b)/2b$, and its variance is $2b^2$. The \mathcal{Q} -function is defined as $\mathcal{Q}(x) := \frac{1}{\sqrt{2\pi}} \int_x^{\infty} e^{-\frac{u^2}{2}} du$. Now for $\epsilon > 0$, $0.5 \geq \delta \geq 0$, let $K = \mathcal{Q}^{-1}(\delta)$ and define $\kappa_{\delta, \epsilon} = \frac{1}{2\epsilon}(K + \sqrt{K^2 + 2\epsilon})$, which can be shown to behave roughly as $O(\ln(1/\delta))^{1/2}/\epsilon$.

Theorem 1: Let G be a system with m inputs and n outputs. Then the mechanism $M(u) = Gu + w$, where all $w_{k,i}, k \geq 0, 1 \leq i \leq n$, are independent Laplace random variables with parameter $b = (\Delta_1 G)/\epsilon$, is ϵ -differentially private for Adj . If w_k is instead a white Gaussian noise with covariance matrix $\kappa_{\delta, \epsilon}^2 (\Delta_2 G)^2 I_n$, the mechanism is (ϵ, δ) -differentially private.

D. Input and Output Perturbation

We see that the amount of noise necessary for differential privacy with the mechanisms of Theorem 1 is proportional to $\Delta_1 G/\epsilon$ or to $\kappa_{\delta, \epsilon} \Delta_2 G$. A very useful additional result stated here informally says that post-processing a differentially private signal without re-accessing the privacy-sensitive input signal does not change the differential privacy guarantee [9, Theorem 1]. Now in Theorem 1 the system G can simply be the identity, whose ℓ_1 - and ℓ_2 - sensitivity for the adjacency relation (4) when $|\cdot|_{\mathcal{Y}}$ is the 1-norm or the 2-norm are $K/(1 - \alpha)$ and $K/\sqrt{1 - \alpha^2}$ respectively. This immediately gives a first possible design for our privacy-preserving observer, simply adding Laplace or Gaussian noise directly to the input signal y , see Fig. 1 a). Moreover the observer can then be designed to mitigate the effect of this input noise, whose distribution is known. We call this design an input perturbation mechanism. Note also that for α close to 1, $\frac{1}{\sqrt{1 - \alpha^2}}$ can be significantly smaller than $\frac{1}{1 - \alpha}$,

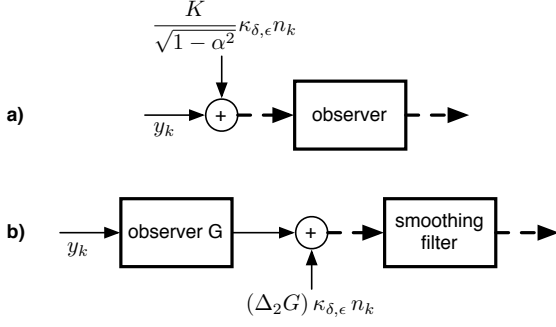


Fig. 1. Gaussian mechanisms with input (a) and output (b) perturbation. n_k represents a zero-mean white Gaussian noise with identity covariance matrix. Dashed lines represent a differentially private signal.

so that sacrificing some δ in the privacy guarantee to use the ℓ_2 -sensitivity can provide better accuracy.

The simple input perturbation mechanism is attractive and can perform well. However, it can also potentially exhibit the following drawbacks. First, the convergence of nonlinear observers is often local and adding noise at the input can lead to poor performance and perhaps divergence of the estimate from the true trajectory. Second, characterizing the output error due to the privacy-preserving noise requires understanding how this noise is transformed after passing through the nonlinear observer. In general, at the output the noise distribution can become multimodal and the noise non white and non zero mean, creating in particular a systematic bias that can be hard to predict. An alternative is the output perturbation mechanism, shown on Fig. 1 b). In this case the privacy-preserving noise is added after the observer denoted G , which from Theorem 1 requires computing the sensitivity of G . In this case we should try to design an observer that has both good tracking performance for the state trajectory and low sensitivity to reduce the output noise necessary, and we focus on this issue in the following. As shown on Fig. 1 b), we can also add a smoothing filter at the output to filter out the Laplace or Gaussian noise, although this will generally affect some transient performance measure of the overall system. We do not discuss the design of the smoothing filter in this paper.

Example 1: Consider the memoryless system $y_k \mapsto \phi(y_k) = y_k^2$ and the adjacency relation (4) for $\alpha = 0$, so that we have a deviation at a single time period of at most K between y_k and \tilde{y}_k . Consider then the Gaussian mechanism, and let's assume $\kappa_{\delta, \epsilon} = 1$. For the input perturbation scheme, the signal $z_k = (y_k + K\xi_k)^2 = y_k^2 + 2Ky_k\xi_k + K^2\xi_k^2 = y_k^2 + e_k$, is differentially private when ξ_k is a standard Gaussian white noise. In this case, the privacy-preserving noise at the input induces a systematic bias at the output between z_k and y_k^2 equal to $\mathbb{E}[e_k] = \mathbb{E}[K^2\xi_k^2] = K^2$.

III. CONTRACTING SYSTEMS

In the rest of the paper we focus on output perturbation mechanisms, as described on Fig. 1 b), and we use contraction theory to bound the sensitivity ΔG and hence compute

the noise level necessary for privacy. Contraction theory has seen significant developments in the past two decades, see, e.g., [11]–[14] and the references therein for references to earlier work. In this section, we present some results that we rely on later in the paper. Proofs of these results are given for completeness, since most results in this area are typically stated for continuous-time rather than discrete-time systems.

Consider a discrete-time system

$$x_{k+1} = f_k(x_k), \quad (6)$$

with $x_k \in X$, for all $k \in \mathbb{N}$. Let us denote by $\phi(k; k_0, x_0)$ the value at time $k \geq k_0$ of a solution of (6) which takes the value x_0 at time k_0 . A forward invariant set for the system (6) is a set $C \subset \mathbb{R}^n$ such that if $x_0 \in C$, then for all k_0 and all $k \geq k_0$, $\phi(k; k_0, x_0) \in C$.

Definition 3: Let α be a nonnegative constant. The system (6) is said to be α -contracting for the norm $|\cdot|_X$ on a forward invariant set $C \subset X$ if for any $k_0 \in \mathbb{N}$ and any two initial conditions $x_0, \tilde{x}_0 \in C$, we have, for all $k \geq k_0$,

$$|\phi(k; k_0, x_0) - \phi(k; k_0, \tilde{x}_0)|_X \leq \alpha^{k-k_0} |x_0 - \tilde{x}_0|_X. \quad (7)$$

Theorem 2: A sufficient condition for the system (6) to be α -contracting for a norm $|\cdot|_X$ on a convex forward invariant set C is that

$$\|F_k(x)\|_X \leq \alpha, \quad \forall x \in C, \forall k \in \mathbb{N}, \quad (8)$$

where $F_k(x) = \frac{\partial f_k}{\partial x}(x)$ is the Jacobian matrix of f_k at x and $\|\cdot\|_X$ is the matrix norm induced by $|\cdot|_X$.

Proof: Consider the path $\gamma(r) = x_0 + r(\tilde{x}_0 - x_0)$, for $r \in [0, 1]$, between the initial conditions x_0 and \tilde{x}_0 . This path is transported into the sequence of functions $\psi_k(r) := \phi(k; k_0, \gamma(r))$. Now define the tangent vectors $w_k(r) := \frac{d}{dr}\psi_k(r) = \psi'_k(r)$. We obtain immediately

$$\begin{aligned} w_{k+1}(r) &= \frac{d\psi_{k+1}(r)}{dr} = \frac{d}{dr}f_k(\psi_k(r)) = \frac{\partial f_k}{\partial x}(\psi_k(r))\psi'_k(r) \\ w_{k+1}(r) &= F_k(\psi_k(r))w_k(r), \quad \forall r \in [0, 1], \forall k \geq k_0. \end{aligned}$$

Then, with $x_k = \phi(k; k_0, x_0)$ and $\tilde{x}_k = \phi(k; k_0, \tilde{x}_0)$,

$$\begin{aligned} |\tilde{x}_k - x_k| &\leq |\psi_k(1) - \psi_k(0)| = \left| \int_0^1 \psi'_k(r) dr \right| \\ &\leq \int_0^1 |w_k(r)| dr \leq \alpha^{k-k_0} \int_0^1 |w_0(r)| dr \\ &= \alpha^{k-k_0} \int_0^1 |\gamma'(r)| dr = \alpha^{k-k_0} |\tilde{x}_0 - x_0|. \end{aligned}$$

For any positive definite matrix P , $|x|_P = \sqrt{x^T P x}$ defines a norm on $X = \mathbb{R}^n$. Specializing the condition of Theorem 2 to this norm, we obtain the following result.

Corollary 1: Let P be a positive definite matrix. A sufficient condition for the system (6) to be α -contracting for the norm $|\cdot|_P$ on a convex forward invariant set C is that the following Linear Matrix Inequalities (LMI) are satisfied

$$F_k(x)^T P F_k(x) \preceq \alpha P, \quad \forall x \in C, \forall k \in \mathbb{N}.$$

Proof: Condition (8) for the matrix norm induced by $|\cdot|_P$ can be rewritten $\|DF_k(x)D^{-1}\|_2 \leq \alpha$, $\forall x, \forall k \in \mathbb{N}$,

where $\|A\|_2$ denotes the induced 2-norm of the matrix A , i.e., its largest singular value, and D is the positive-definite square root of P . The equivalence with the LMI is immediate. ■

Remark 2: Contraction theory can be developed in a more general differential geometric framework [11], [13], which we do not use here however, for simplicity of exposition and also because some of the needed explicit calculations become more difficult, e.g., requiring the computation of non-trivial geodesic paths and distances.

Under conditions such as that of Theorem 2, cascades of contracting systems are again contracting [11], [12]. Consider the system (6) on $X = \mathbb{R}^n$ equipped with the norm $|\cdot|_X$, and assumes that it satisfies condition (8). Then, consider another system $z_{k+1} = g_k(x_k, z_k)$, with $z_k \in Z = \mathbb{R}^{n'}$ equipped with the norm $|\cdot|_Z$, and assume that we have the bounds

$$\|G_k(x, z)\|_Z \leq \beta, \quad \forall x \in C, \forall z \in C', \forall k \in \mathbb{N}, \quad (9)$$

$$\|A_k(x, z)\|_{X,Z} \leq K, \quad \forall x \in C, \forall z \in C', \forall k \in \mathbb{N}, \quad (10)$$

where $G_k(x, z) = \frac{\partial g_k}{\partial z}(x, z)$, $A_k(x, z) = \frac{\partial g_k}{\partial x}(x, z)$, β, K are nonnegative constants, C' is convex and $C \times C'$ is forward invariant for the coupled system.

Theorem 3: Under the previous conditions (8), (9), (10), for any $\rho > 0$ the cascade system

$$\begin{cases} x_{k+1} = f_k(x_k) \\ z_{k+1} = g_k(x_k, z_k) \end{cases}$$

is γ -contracting on $X \times Z$ for the norm

$$|x^T y^T|^T = \rho|x|_X + |z|_Z, \quad (11)$$

with $\gamma = \max \left\{ \alpha + \frac{1}{\rho}K, \beta \right\}$. More precisely, the Jacobian of the cascade system $J_k(x, z) = \begin{bmatrix} F_k(x) & 0 \\ A_k(x, z) & G_k(x, z) \end{bmatrix}$ satisfies

$$\|J_{k+1}(x, z)\| \leq \gamma, \quad \forall x \in C, z \in C', \forall k \in \mathbb{N}, \quad (12)$$

where $\|\cdot\|$ is the matrix norm induced by the norm (11).

Proof: Let $(v, w) \in \mathbb{R}^n \times \mathbb{R}^{n'}$. Then

$$\begin{aligned} \left| J_{k+1}(x, z) \begin{bmatrix} v \\ w \end{bmatrix} \right| &= \rho|F_k(x)v|_X + |A_k(x, z)v + B_k(x, z)w|_Z \\ &\leq \alpha\rho|v|_X + K|v|_X + \beta|w|_Z \\ &= \rho(\alpha + K/\rho)|v|_X + \beta|w|_Z \\ &\leq \gamma(\rho|v|_X + |w|_Z) = \gamma|v^T w^T|^T, \end{aligned}$$

which proves (12). ■

Note that in Theorem 3 we need to choose ρ large enough to satisfy the condition $\alpha + \frac{1}{\rho}K < 1$ to show that pairwise trajectories of the cascade system are effectively converging toward each other. We can now prove the following result, which will be our main tool in the following.

Theorem 4: Consider a (contracting) system on X

$$\bar{x}_{k+1} = f_k(\bar{x}_k), \quad (13)$$

and the modified system

$$x_{k+1} = f_k(x_k) + d_k(x_k), \quad (14)$$

where $d_k(x_k)$ denotes a perturbation input. Suppose that there exists $k_0 \in \mathbb{N}$ such that $d_k(x_k) = 0$ for $k < k_0$, and

$$|d_k(x_k)|_X \leq K\alpha^{k-k_0}, \quad \forall k \geq k_0, \quad (15)$$

for some constants $K, \alpha \geq 0$. Finally, suppose that we have the contraction condition

$$\|J_k(x; p)\|_X \leq \beta, \quad \forall p \in [0, 1], \forall x \in C, \forall k \geq k_0, \quad (16)$$

where C is a convex set that is forward invariant for (13) and (14), and

$$J_k(x; p) = \frac{\partial f_k}{\partial x}(x) + \frac{p}{\gamma^{k-k_0}} \frac{\partial d_k}{\partial x}(x).$$

If $x_0, \bar{x}_0 \in C$, then for $k \geq k_0$, and any $\rho > 0$, we have

$$|x_k - \bar{x}_k|_X \leq \rho(\gamma^{k-k_0} - \alpha^{k-k_0}) + \gamma^{k-k_0}|x_0 - \bar{x}_0|_X,$$

where $x_k = \phi(k; k_0, x_0)$, $\bar{x}_k = \phi(k; k_0, \bar{x}_0)$ and $\gamma = \max \left\{ \alpha + \frac{K}{\rho}, \beta \right\}$.

Proof: Following the idea in [12, Lemma 4] for example, we consider the following cascade system with $p_k \in [0, 1]$

$$\begin{aligned} p_{k+1} &= \alpha p_k \\ x_{k+1} &= f_k(x_k) + \frac{p}{\gamma^{k-k_0}} d_k(x_k). \end{aligned}$$

For the initial condition $(0, \bar{x}_0)$ at k_0 , we obtain a trajectory of the unperturbed system (13), whereas for the initial condition $(1, x_0)$, we obtain a trajectory of the perturbed system (14). The scalar p system is α -contracting. For each $p \in [0, 1]$, the x -system is β -contracting by (16). Moreover, the differential of the second vector field with respect to p is $d_k(x)/\gamma^{k-k_0}$, which is bounded by K from (15). Hence, applying the result of Theorem 3, for any $\rho > 0$ the overall system is contracting with respect to the norm $\rho|p| + |x|$ (where $p \in \mathbb{R}, x \in \mathbb{R}^n$), with rate $\gamma = \max \left\{ \alpha + \frac{K}{\rho}, \beta \right\}$, so

$$\begin{aligned} \rho\alpha^{k-k_0} + |x_k - \bar{x}_k|_X &\leq \gamma^{k-k_0}(\rho + |x_0 - \bar{x}_0|_X) \\ |x_k - \bar{x}_k|_X &\leq \rho(\gamma^{k-k_0} - \alpha^{k-k_0}) + \gamma^{k-k_0}|x_0 - \bar{x}_0|_X. \end{aligned}$$

■
Remark 3: Note that if d_k is independent of x , then the contraction condition (16) is simply a contraction condition on the original system (13) since $\frac{\partial d_k(x)}{\partial x} = 0$.

IV. DIFFERENTIALLY PRIVATE OBSERVERS WITH OUTPUT PERTURBATION

Let us now return to our initial differentially private observer design problem with output perturbation. We can rewrite the system (3) in the form $z_{k+1} = (f_k(z_k) - L_k g_k(z_k)) + L_k y_k$. For a measured signal \tilde{y} adjacent to y according to (4), we then get the observer state trajectory

$$\begin{aligned} \tilde{z}_{k+1} &= (f_k(\tilde{z}_k) - L_k g_k(\tilde{z}_k)) + L_k \tilde{y}_k \\ \tilde{z}_{k+1} &= (f_k(\tilde{z}_k) - L_k g_k(\tilde{z}_k)) + L_k y_k + L_k \delta_k, \end{aligned} \quad (17)$$

where $\delta_k = \tilde{y}_k - y_k$. We can now use the gain matrices L_k to attempt to design a contractive observer (in order for z_k to converge to x_k), while at the same time minimizing

the “gain” of the map $\delta \rightarrow z$. The proof of the following proposition follows immediately from Theorem 4.

Proposition 1: Consider the system (3), and two measured signals y, \tilde{y} adjacent according to (4). Let $K' = K \times \sup_k \|L_k\|_{X,Y}$. Suppose also that we have the bound

$$\|F_k(z) - L_k G_k(z)\|_X \leq \beta, \quad \forall z \in C, \forall k \in \mathbb{N}, \quad (18)$$

for some constant β , where $F_k(z) = \frac{\partial f_k}{\partial z}(z)$, $G_k(z) = \frac{\partial g_k}{\partial z}(z)$, and $C \subset X$ is a convex forward invariant set for (3) and (17). Then for the two trajectories z_k and \tilde{z}_k of (3) corresponding to the inputs y_k and \tilde{y}_k (and assuming the same initial condition $z_0 = \tilde{z}_0 \in C$ for our observer), we have for any $\rho > 0$

$$\begin{cases} z_k = \tilde{z}_k, & \forall k \leq k_0 \\ \|z_k - \tilde{z}_k\|_X \leq \rho(\gamma^{k-k_0} - \alpha^{k-k_0}), & \forall k > k_0, \end{cases}$$

where $\gamma = \max\left\{\alpha + \frac{K'}{\rho}, \beta\right\}$, and k_0 is the time period where y and \tilde{y} start to potentially differ according to (4).

Note in the previous proposition that the choice of L_k has an impact both on ρ and γ . Increasing the gain L_k can help decrease the contraction rate β , but at the same time it increases K' , forcing us to increase ρ so that $\alpha + K'/\rho < 1$. Hence in general we should look to achieve a reasonable contraction rate β with the smallest gain possible, in order to reduce the overall system sensitivity (in the sense of Section II-C). We conclude this section with two corollaries of Proposition 1 providing differentially private observers with output perturbation.

Corollary 2: Consider the signal $\hat{x}_k = z_k + \xi_k$, where z_k is computed from (3), the conditions of Proposition 1 are satisfied for the 1-norm on X , and $\xi_{k,i}$ are iid Laplace random variables with parameter

$$b = \frac{\rho}{\epsilon} \left(\frac{1}{1-\gamma} - \frac{1}{1-\alpha} \right). \quad (19)$$

Then this signal \hat{x}_k is ϵ -differentially private for the adjacency relation (4).

Corollary 3: Let P be a positive definite matrix. Consider the signal $\hat{x}_k = z_k + \xi_k$, where z_k is computed from (3), the conditions of Proposition 1 are satisfied for the $\|\cdot\|_P$ norm on X , and ξ_k is a Gaussian white noise with covariance matrix $\sigma^2 P^{-1}$, where $\sigma = \kappa_{\delta,\epsilon} \rho B$ and

$$B := \left(\sum_{k \geq 0} (\gamma^k - \alpha^k)^2 \right)^{1/2} \leq \frac{1}{\sqrt{1-\gamma^2}}.$$

Then this signal \hat{x}_k is (ϵ, δ) -differentially private for the adjacency relation (4).

Proof: From the bound of Proposition 1, we deduce that $Dz_k + \zeta_k$ is a differentially private signal, where ζ_k is a Gaussian white noise with covariance matrix $\sigma^2 I$ and D is the matrix square root of P . Hence $D^{-1}(Dz_k + \zeta_k)$ is also differentially private and we defined $\xi_k = D^{-1}\zeta_k$. ■

We thus have two differentially private mechanisms with output perturbation, provided we can design the matrices L_k

to verify the assumptions of Proposition 1 with the 1- or 2-norm on X . The next sections provide application examples for the methodology.

V. EXAMPLE I: ESTIMATING LINK FORMATION PREFERENCES IN DYNAMIC SOCIAL NETWORKS

Statistical studies of networks have intensified tremendously in recent years, with one motivating application being the emergence of online social networking communities. In this section we focus on a state-space model recently proposed in [15] to describe the dynamics of link formation in networks, called the Dynamic Stochastic Blockmodel. This model combines a linear state-space model for the underlying dynamics of the network and the stochastic blockmodel of Holland et al. [16], resulting in a nonlinear measurement equation. Examples of applications of this model include mining email and cell phone databases [15], which obviously contain privacy-sensitive data.

Consider a set of n nodes. Each node corresponds to an individual and can belong to one of N classes. Let θ_k^{ab} be the probability of forming an edge at time k between a node in class a and a node in class b , and let θ_k denote the vector of probabilities $[\theta_k^{ab}]_{1 \leq a, b \leq N}$. For example, edges could represent email exchanges or phone conversations. Edges are assumed to be formed independently of each other according to θ_k . Let $y_k^{ab} = \frac{m_k^{ab}}{n^{ab}}$ be the observed density of edges between classes a and b , where m_k^{ab} is the number of observed edges between classes a and b at time k , and n^{ab} is the maximum possible number of edges between these two classes. For simplicity, we assume that the quantities n^{ab} are publicly known (for example, if the class of each node is public information), and we focus on the problem of estimating the parameters θ_k^{ab} using the signals y_k^{ab} . This corresponds to the “a priori” blockmodeling setting in [15], [16]. The links formed between specific nodes constitute private information however, so directly releasing m_k^{ab} or y_k^{ab} or an estimate based on them is not allowed.

If n^{ab} is large enough, the authors in [15] argue from the Central Limit Theorem that an approximate model where y_k^{ab} is Gaussian is justified, so that

$$y_k = \theta_k + v_k, \quad (20)$$

where v_k is a Gaussian noise vector with diagonal covariance matrix V_k (whose entries theoretically should depend on θ_k , but this aspect is neglected in the model). Rather than defining a dynamic model for θ_k , whose entries are constrained to be between 0 and 1, let us redefine the state vector to be the so-called logit of θ_k , denoted ψ_k , with entries $\psi_k^{ab} = \ln \frac{\theta_k^{ab}}{1-\theta_k^{ab}}$, which are well defined for $0 < \theta_k^{ab} < 1$. The dynamics of ψ_k is assumed to be linear

$$\psi_{k+1} = F\psi_k + w_k, \quad (21)$$

for some known matrix F . The noise vectors w_k are assumed to be iid Gaussian with known covariance matrix W in [15]. The observation model (20) now becomes

$$y_k = g(\psi_k) + v_k, \quad (22)$$

where the components of g are given by the logistic function applied to each entry of ψ , i.e.,

$$g^{ab}(\psi_k) = \frac{1}{(1 + e^{-\psi_k^{ab}})}.$$

An Extended Kalman Filter (EKF) is proposed in [15] to estimate ψ , but we pursue here a deterministic observer design to illustrate the ideas discussed in the previous sections. Hence, we consider an observer of the form

$$\hat{\psi}_{k+1} = F_k \hat{\psi}_k + L(y_k - g(\hat{\psi}_k)) = (F_k \hat{\psi}_k - Lg(\hat{\psi}_k)) + Ly_k,$$

with L a constant square gain matrix. To enforce contraction as in Proposition 1, we should choose L so that $\|F_k - LG(\psi_k)\| \leq \beta$, where $G(\psi)$ is the Jacobian of g at ψ , a square and diagonal matrix with entries $G^{ii}(\psi) = \frac{e^{-\psi^i}}{(1+e^{-\psi^i})^2}$, with i indexing the pairs (a, b) . The only non-linearity in the model (21), (22) comes from the observation model (22).

To simplify the following discussion, let's assume that F is also diagonal (as in [15], where the coupling between components occurs only through the non-diagonal covariance matrix W). In this case, the systems completely decouple into scalar systems, and it is natural to choose L to be diagonal as well. The observer for one of these scalar system takes the form

$$z_{k+1} = fz_k + l \left(y_k - \frac{1}{1 + e^{-z_k}} \right) = fz_k - \frac{l}{1 + e^{-z_k}} + ly_k, \quad (23)$$

where z_k is one component (a, b) of $\hat{\psi}_k$ and y_k now represents just the corresponding scalar component of the measurement vector as well. Since the state space \mathbf{X} is now \mathbb{R} , the norm $\|\cdot\|_X$ is simply the absolute value. For contraction, we wish to impose the condition, for some $0 < \beta < 1$,

$$-\beta \leq f - \frac{le^{-z}}{(1 + e^{-z})^2} \leq \beta \quad (24)$$

$$\text{i.e., } f - \beta \leq \frac{le^{-z}}{(1 + e^{-z})^2} \leq f + \beta. \quad (25)$$

Now note that $0 \leq \frac{e^{-z}}{(1+e^{-z})^2} \leq \frac{1}{4}$ for all z . Hence, by taking $l \leq 4(f + \beta)$, the right hand side of (25) is satisfied. Moreover, for $-a \leq z \leq a$, we have $\frac{e^{-z}}{(1+e^{-z})^2} \geq b := \frac{e^{-a}}{(1+e^{-a})^2}$. In this case, by taking $l \geq \frac{f-\beta}{b}$, the left hand side of (25) is also satisfied.

Suppose that we want to design a privacy-preserving observer for the interval $\theta \in [0.05, 0.95]$, or equivalently $\psi \in [-2.95, 2.95]$ approximately. In this interval, we have

$$0.0475 \leq \frac{e^{-\psi}}{(1 + e^{-\psi})^2} \leq \frac{1}{4}.$$

Suppose that we have $f = 0.95$. Then we must have

$$\frac{f - \beta}{0.0475} \leq l \leq 4(f + \beta). \quad (26)$$

In general to reduce the sensitivity we should choose a small gain l , which is compatible with (26) if we choose β close enough to f . Indeed, setting $l = (f - \beta)/0.0475$ and $\rho =$




Fig. 2. Estimate of the edge formation probability θ_k^{ab} , for some classes (a, b) . The measured edge density is generated from one component of the model (20), (21) with $f = 0.95$ and w_k, v_k iid Gaussian random variables with zero mean and standard deviation 0.05 and 0.01 respectively. The gain l of the observer (23) is set to 0.3. We plot $1/(1 + \exp(-z_k))$ as our estimate of θ_k , where z_k is a 1-differentially private estimate of ψ_k with no postfiltering, for the adjacency relation (4) with parameter values detailed in the main text.

$lM/(\beta - \alpha)$ in Proposition 1 so that $\gamma = \beta$ (assuming $\beta > \alpha$), we can verify that the ℓ_1 sensitivity say and thus the noise parameter b in (19) decreases monotonically as β increases toward f . However, performance concerns for the observer should also dictate the minimum tolerable gain (with a gain $l = 0$, the observer is perfectly private but is not useful).

Suppose the disturbance tolerated by the adjacency relation satisfies the bound (4) with $K = 10^{-3}$ and $\alpha = 0.25$. That is, for the pair of classes (a, b) under consideration, we want to provide a differential privacy guarantee making it hard to detect a transient variation in the number of created edges, as long as this variation represents initially at most 0.1% of all the edges between classes a and b , and subsequently decreases geometrically at rate $1/4$. Concretely if edges represent phone conversations for example, this means that if an individual in class a suddenly increases his call volume with class b but by an amount representing less than 0.1% of all calls between a and b , and then reduces this temporary activity at rate α , detection of this event by any means from a differentially private estimate of ψ_k^{ab} will necessarily have a low probability of success. If a gain $l = 0.3$ say is judged to be still adequate for the application in terms of tracking performance, we can take $\beta = f - 0.0475l \approx 0.936$ and we get $b = 6.23 \times 10^{-3}/\epsilon$ in (19). If we publish $z_k + \xi_k$ with ξ_k a Laplace white noise with this parameter b , we obtain an ϵ -differentially private estimator of ψ_k . Figure 2 illustrates the behavior of the resulting privacy-preserving observer.

VI. EXAMPLE II: SYNDROMIC SURVEILLANCE

Syndromic surveillance systems monitor health related data in real-time in a population to facilitate early detection of epidemic outbreaks [17]. In particular, recent studies have shown the correlation between certain non-medical data, e.g., search engine queries related to a specific disease, and the proportion of individuals infected by this disease in the population [18]. Although time series analysis can be used

to detect abnormal patterns in the collected data [17], here we focus on a model-based filtering approach [19], and develop a differentially private observer for a 2-dimensional epidemiological model.

The following SIR model of Kermack and McKendrick [20], [21] models the evolution of an epidemic in a population by dividing individuals into 3 categories: susceptible (S), i.e., individuals who might become infected if exposed; infectious (I), i.e., currently infected individuals who can transmit the infection; and recovered (R) individuals, who are immune to the infection. A simple version of the model in continuous-time includes bilinear terms and reads

$$\begin{aligned}\frac{ds}{dt} &= -\mu\mathcal{R}_o i s \\ \frac{di}{dt} &= \mu\mathcal{R}_o i s - \mu i.\end{aligned}$$

Here i and s represent the proportion of the total population in the classes I and S . The last class R need not be included in this model because we have the constraint $i + s + r = 1$. The parameter \mathcal{R}_o is called the basic reproduction number and represents the average number of individuals infected by a sick person. The epidemic can propagate when $\mathcal{R}_o > 1$. The parameter μ represents the rate at which infectious people recover and move to the class R . More details about this model can be found in [21].

Discretizing this model with sampling period τ , we get the discrete-time model

$$s_{k+1} = s_k - \tau\mu\mathcal{R}_o i_k s_k + w_{1,k} = f_1(s_k, i_k) + w_{1,k} \quad (27)$$

$$i_{k+1} = i_k + \tau\mu i_k (\mathcal{R}_o s_k - 1) + w_{2,k} = f_2(s_k, i_k) + w_{2,k}, \quad (28)$$

where we have also introduced noise signals w_1 and w_2 in the dynamics. We assume here for simplicity that we can collect syndromic data providing a noisy measurement of the proportion of infected individuals, i.e.,

$$y_k = i_k + v_k,$$

where v_k is a noise signal. We can then consider the design of an observer of the form

$$\begin{aligned}\hat{s}_{k+1} &= f_1(\hat{s}_k, \hat{i}_k) + l_1(y_k - \hat{i}_k) \\ \hat{i}_{k+1} &= f_2(\hat{s}_k, \hat{i}_k) + l_2(y_k - \hat{i}_k).\end{aligned}$$

We define the Jacobian matrix of the system (27), (28)

$$J(s, i) = I_2 + \tau\gamma\mathcal{R}_o \begin{bmatrix} -i & -s \\ i & s - 1/\mathcal{R}_o \end{bmatrix},$$

as well as the gain matrix $L = [l_1, l_2]^T$ and observation matrix $C = [0, 1]$.

Following Corollary 3 and according to Corollary 1, the contraction rate constraint (18) for a norm $|\cdot|_P$ on \mathbb{R}^2 with P a positive definite matrix is equivalent to the family of inequalities

$$\begin{aligned}(J(s, i) - LC)^T P (J(s, i) - LC) &\preceq \beta P \\ J_x^T P J_x - J_x^T P L C - C^T L^T P J_x + C^T L^T P L C &\preceq \beta P,\end{aligned}$$

where we used $J_x := J(s, i)$ to simplify the notation. Defining the new variable $X = PL$, this can be rewritten

$$J_x^T P J_x - J_x^T X C - C^T X^T J_x + C^T X^T P^{-1} X C \preceq \beta P,$$

which, using the Schur complement, is equivalent to the family of LMIs

$$\begin{bmatrix} \beta P - J_x^T P J_x + J_x^T X C + C^T X^T J_x & C^T X^T \\ X C & P \end{bmatrix} \succeq 0, \quad (29)$$

for all $x = (s, i)$ in the region where we want to prove contraction. If we can find P, X satisfying these inequalities, we recover the observer gain matrix simply as $L = P^{-1}X$.

Note that to minimize K' in Proposition 1, we should try to minimize $\|L\|_P^2 = L^T P L = X^T P^{-1} X$, or equivalently minimize g_1 such that the following LMI is satisfied

$$\begin{bmatrix} g_1 & X^T \\ X & P \end{bmatrix} \succeq 0. \quad (30)$$

However, we should also minimize P^{-1} , which appears in the covariance matrix of the privacy-preserving noise in Corollary 3, or equivalently minimize g_2 subject to

$$\begin{bmatrix} g_2 I & I \\ I & P \end{bmatrix} \succeq 0. \quad (31)$$

In the end, we choose to minimize a cost function of the form $g_1 + c g_2$, with c a coefficient appropriately tuned to balance observer gain and level of privacy-preserving noise, subject to the LMI constraints (29), (30) and (31), and $P \succ 0$ or perhaps $P \succeq c' I$ for another constant c' if we wish to impose a hard upper bound on the noise covariance.

Example 2: Let's assume $\mu = 0.1$, $\mathcal{R}_o = 3$, $M = 5 \times 10^{-4}$, $\alpha = 0.25$ in (4), and $\epsilon = 2$, $\delta = 0.05$. That is, we wish to provide a $(2, 0.05)$ -differential privacy guarantee for maximum deviations of 0.05% (see the discussion in the previous section). Although not a perfectly rigorous contraction certificate, we sample the continuous set of constraints (29) by sampling the set $\{(s, i) | 0.01 \leq i \leq 0.5, 0 \leq s \leq 1 - i\}$ at the values of s, i multiple of 0.01, to obtain a finite number of LMIs. A more rigorous approach to enforce these constraints could make use of sum-of-squares programming [22]. Following the procedure above, for the choice $\beta = 1 - 10^{-5}$, $c = 1$, we obtain the observer gain $L = [-0.3657; 0.2951]$ and the covariance matrix

$$\sigma^2 P^{-1} = \begin{bmatrix} 0.3 & -0.11 \\ -0.11 & 0.13 \end{bmatrix} \times 10^{-4}$$

for the Gaussian privacy-preserving noise. A typical sample trajectory of the estimate of i is shown on Fig. 3.

VII. CONCLUSION

We have discussed input and output perturbation mechanisms to design model-based nonlinear estimators with differential privacy guarantees. In general, we wish to achieve a good contraction rate with the smallest gain possible, and in fact this idea applies to both types of mechanisms. Future work includes comparing quantitatively input and output perturbation schemes, and generalizing both by combining




Fig. 3. Estimate of the number of infectious people over time produced by the observer. The noise standard deviations were set to $\sigma_{v_k} = 0.02$ and $\sigma_{w_k} = 0.01\tau$ respectively. The output of the privacy-preserving observer is not filtered.

pre- and post-processing as illustrated on Fig. 1 b) and in [9], [10] for the linear case.

REFERENCES

- [1] A. Narayanan and V. Shmatikov, "Robust de-anonymization of large sparse datasets (how to break anonymity of the Netflix Prize dataset)," in *Proceedings of the IEEE Symposium on Security and Privacy*, 2008.
- [2] J. A. Calandrino, A. Kilzer, A. Narayanan, E. W. Felten, and V. Shmatikov, "'you might also like': Privacy risks of collaborative filtering," in *Proceedings of the IEEE Symposium on Security and Privacy*, Berkeley, CA, May 2011.
- [3] D. H. Wilson and C. Atkeson, "Simultaneous tracking and activity recognition (STAR) using many anonymous, binary sensors," in *Pervasive Computing*, ser. Lecture Notes in Computer Science, H.-W. Gellersen, R. Want, and A. Schmidt, Eds. Springer Berlin Heidelberg, 2005, vol. 3468, pp. 62–79.
- [4] L. Sankar, W. Trappe, K. Ramchandran, H. V. Poor, and M. Debbah, Eds., *IEEE Signal Processing Magazine, Special issue on Signal Processing for Cybersecurity and Privacy*, September 2013.
- [5] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Proceedings of the Third Theory of Cryptography Conference*, 2006, pp. 265–284.
- [6] C. Dwork, M. Naor, T. Pitassi, and G. N. Rothblum, "Differential privacy under continual observations," in *Proceedings of the ACM Symposium on the Theory of Computing (STOC)*, Cambridge, MA, June 2010.
- [7] T.-H. H. Chan, E. Shi, and D. Song, "Private and continual release of statistics," *ACM Transactions on Information and System Security*, vol. 14, no. 3, pp. 26:1–26:24, November 2011.
- [8] J. Le Ny and G. J. Pappas, "Differentially private Kalman filtering," in *Proceedings of the 50th Annual Allerton Conference on Communication, Control, and Computing*, October 2012.
- [9] —, "Differentially private filtering," *IEEE Transactions on Automatic Control*, vol. 59, no. 2, pp. 341–354, February 2014.
- [10] J. Le Ny and M. Mohammady, "Differentially private MIMO filtering for event streams and spatio-temporal monitoring," in *IEEE Conference on Decision and Control*, Los Angeles, CA, December 2014.
- [11] W. Lohmiller and J.-J. Slotine, "On contraction analysis for non-linear systems," *Automatica*, vol. 34, no. 6, pp. 683–696, 1998.
- [12] E. D. Sontag, "Contractive systems with inputs," in *Perspectives in Mathematical System Theory, Control, and Signal Processing*, J. Willems, S. Hara, Y. Ohta, and H. Fujioka, Eds. Springer-Verlag, 2010, pp. 217–228.
- [13] F. Forni and R. Sepulchre, "A differential Lyapunov framework for contraction analysis," *IEEE Transactions on Automatic Control*, vol. 59, no. 3, pp. 614–628, March 2014.
- [14] D. Angeli, "A Lyapunov approach to incremental stability properties," *IEEE Transactions on Automatic Control*, vol. 47, no. 3, pp. 410–421, March 2000.
- [15] K. S. Xu and A. O. Hero III, "Dynamic stochastic blockmodels for time-evolving social networks," *Journal of Selected Topics in Signal Processing*, vol. 8, no. 4, pp. 552–562, August 2014, Special Issue on Signal Processing for Social Networks.
- [16] P. W. Holland, K. B. Laskey, and S. Leinhardt, "Stochastic blockmodels: First steps," *Social Networks*, vol. 5, no. 2, pp. 109–137, 1983.
- [17] A. B. Lawson and K. Kleinman, *Spatial and Syndromic Surveillance for Public Health*. Wiley, 2005.
- [18] J. Ginsberg, M. H. Mohebbi, R. S. Patel, L. Brammer, M. S. Smolinski, and L. Brilliant, "Detecting influenza epidemics using search engine query data," *Nature*, vol. 457, pp. 1012–1014, 2009.
- [19] A. Skvortsov and B. Ristic, "Monitoring and prediction of an epidemic outbreak using syndromic observations," *Mathematical biosciences*, vol. 240, pp. 12–19, 2012.
- [20] W. O. Kermack and A. G. McKendrick, "A contribution to the mathematical theory of epidemics," *Proceedings of the Royal Society of London Series A*, vol. 115, pp. 700–721, 1927.
- [21] F. Brauer, P. van den Driessche, and J. Wu, Eds., *Mathematical Epidemiology*, ser. Lecture Notes in Mathematics. Berlin: Springer-Verlag, 2008, vol. 1945.
- [22] E. M. Aylward, P. A. Parrilo, and J.-J. E. Slotine, "Stability and robustness analysis of nonlinear systems via contraction metrics and SOS programming," *Automatica*, vol. 44, no. 8, pp. 2163–2170, August 2008.